



USPTO

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☒ The Guide

THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used: **database partition bulletin**

Found 2 of 71 searched out of 71.

Sort results by

[Save results to a Binder](#)[Try an Advanced Search](#)

Display results

[Search Tips](#)[Try this search in The ACM Guide](#)
☒ Open results in a new window

Results 1 - 2 of 2

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [A distributed object-oriented database system supporting shared and private databases](#)



Won Kim, Nat Ballou, Jorge F. Garza, Darrell Woelk

January 1991 **ACM Transactions on Information Systems (TOIS)**, Volume 9 Issue 1**Publisher:** ACM PressFull text available: [pdf\(1.58 MB\)](#)
 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

ORION-2 is a commercially available, federated, object-oriented database management system designed and implemented at MCC. One major architectural innovation in ORION-2 is the coexistence of a shared database and a number of private databases. The shared database is accessible to all authorized users of the system, while each private database is accessible to only the user who owns it. A distributed database system with a shared database and private databases for individual users is a natu ...

Keywords: client-server architecture, federated databases, object-oriented databases

2 [Database research at Columbia University](#)



Shih-Fu Chang, Luis Gravano, Gail E. Kaiser, Kenneth A. Ross, Salvatore J. Stolfo

September 1998 **ACM SIGMOD Record**, Volume 27 Issue 3**Publisher:** ACM PressFull text available: [pdf\(659.46 KB\)](#) Additional Information: [full citation](#), [index terms](#)

Results 1 - 2 of 2

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)
Useful downloads: [Adobe Acrobat](#)[QuickTime](#)[Windows Media Player](#)[Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☒ The Guide



THE ACM DIGITAL LIBRARY


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used: **query partition**

Found 13 of 21

 Sort results by


Save results to a Binder

[Try an Advanced Search](#)

 Display results

[Search Tips](#)
[Try this search in The ACM Guide](#)

[Open results in a new window](#)

Results 1 - 13 of 13

 Relevance scale ☐ ☐ ☐

1 [The complexity of operations on a fragmented relation](#)



C. Meghini, C. Thanos

 March 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 1

Publisher: ACM Press

Full text available: pdf(2.13 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Data fragmentation is an important aspect of distributed database design, in which portions of relations, tailored to the specific needs of local applications, are defined to be further allocated to the sites of the computer network supporting the database system. In this paper we present a theory of fragmentation with overlapping fragments to study the complexity of the problems involved in checking the completeness of a fragmentation schema and in querying and updating a fragmented relation ...

Keywords: NP-hardness, completeness of fragmentation schemas, query optimization, relation fragmentation, updates

2 [Algorithmic program diagnosis](#)



Ehud Y. Shapiro

 January 1982 **Proceedings of the 9th ACM SIGPLAN-SIGACT symposium on Principles of programming languages POPL '82**
Publisher: ACM Press

Full text available: pdf(894.31 KB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The notion of program correctness with respect to an interpretation is defined for a class of programming languages. Under this definition, if a program terminates with an incorrect output then it contains an incorrect procedure. Algorithms for detecting incorrect procedures are developed. These algorithms formalize what experienced programmers may know already. A logic program implementation of these algorithms is described. Its performance suggests that the algorithms can be the backbone of debugging ...

3 [Paper session DB-9 \(databases\): query processing 1: Selectivity-based partitioning: a divide-and-union paradigm for effective query optimization](#)



Neoklis Polyzotis

 October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**

Publisher: ACM Press

Full text available:  pdf(253.90 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern query optimizers select an efficient join ordering for a physical execution plan based essentially on the *average* join selectivity factors among the referenced tables. In this paper, we argue that this "monolithic" approach can miss important opportunities for the effective optimization of relational queries. We propose *selectivity-based partitioning*, a novel optimization paradigm that takes into account the join correlations among *relation fragments* in order to essen ...


4 Two learning schemes in information retrieval



C. T. Yu, H. Mizuno

May 1988 **Proceedings of the 11th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '88**

Publisher: ACM Press

Full text available:  pdf(1.39 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Two methods are given to improve weighting schemes by using relevance information of a set of queries. The first method is to estimate parameter values of two independence models in information retrieval — the binary independence model and the non-binary independence model. The parameters estimated here are used to calculate optimal weights for terms in a different set of queries. Performance of this estimation is compared to the inverse document frequency method, the cosine measure, ...

5 Computation of term associations by a neural network



S. K. M. Wong, Y. J. Cai, Y. Y. Yao

July 1993 **Proceedings of the 16th annual international ACM SIGIR conference on Research and development in information retrieval SIGIR '93**

Publisher: ACM Press

Full text available:  pdf(636.57 KB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Stream processing & query processing II: Towards optimal continuous nearest neighbor queries in spatial databases



Victor Teixeira de Almeida

November 2006 **Proceedings of the 14th annual ACM international symposium on Advances in geographic information systems GIS '06**

Publisher: ACM Press

Full text available:  pdf(451.82 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

One of the most important kind of queries in spatial data-bases to support location-based services (LBS) is the continuous nearest neighbors (CNN) query. Given a spatial data set of points of interest and a moving query point q , the CNN query partitions q into a set of adjacent disjoint intervals associated with their nearest points of interest. Existing solutions to this problem are known to be sub-optimal in terms of disk accesses. In this paper, we present an algorithm to comput ...

Keywords: continuous nearest neighbors, location based services, spatial databases

7 Physical interface: TAG: a Tiny AGgregation service for ad-hoc sensor networks

Samuel Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong
December 2002

ACM SIGOPS Operating Systems Review, Volume 36 Issue S1**Publisher:** ACM Press

Full text available: pdf(2.19 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

We present the Tiny AGgregation (TAG) service for aggregation in low-power, distributed, wireless environments. TAG allows users to express simple, declarative queries and have them distributed and executed efficiently in networks of low-power, wireless sensors. We discuss various generic properties of aggregates, and show how those properties affect the performance of our in network approach. We include a performance study demonstrating the advantages of our approach over traditional centralize ...

8 Advanced SQL modeling in RDBMS

Andrew Witkowski, Srikanth Bellamkonda, Tolga Bozkaya, Nathan Folkert, Abhinav Gupta, John Haydu, Lei Sheng, Sankar Subramanian

March 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 1**Publisher:** ACM Press

Full text available: pdf(279.06 KB)

Additional Information: [full citation](#), [appendices and supplements](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)

Commercial relational database systems lack support for complex business modeling. ANSI SQL cannot treat relations as multidimensional arrays and define multiple, interrelated formulas over them, operations which are needed for business modeling. Relational OLAP (ROLAP) applications have to perform such tasks using joins, SQL Window Functions, complex CASE expressions, and the GROUP BY operator simulating the pivot operation. The designated place in SQL for calculations is the SELECT clause, whi ...

Keywords: Excel, OLAP, analytic computations, spreadsheet**9** Progressive evaluation of nested aggregate queries

Kian-Lee Tan, Cheng Hian Goh, Beng Chin Ooi

December 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**,

Volume 9 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(380.81 KB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

In many decision-making scenarios, decision makers require rapid feedback to their queries, which typically involve aggregates. The traditional *blocking execution model* can no longer meet the demands of these users. One promising approach in the literature, called *online aggregation*, evaluates an aggregation query progressively as follows: as soon as certain data have been evaluated, approximate answers are produced with their respective running confidence intervals; as more data a ...

Keywords: Approximate answers, Multi-threading, Nested aggregate queries, Online aggregation, Progressive query processing**10** Industrial sessions: big data: Automating physical database design in a parallel database

Jun Rao, Chun Zhang, Nimrod Megiddo, Guy Lohman

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02****Publisher:** ACM Press

Full text available: pdf(1.38 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Physical database design is important for query performance in a shared-nothing parallel database system, in which data is horizontally partitioned among multiple independent nodes. We seek to automate the process of data partitioning. Given a workload of SQL statements, we seek to determine automatically how to partition the base data across multiple nodes to achieve overall optimal (or close to optimal) performance for that workload. Previous attempts use heuristic rules to make those decisions ...

11 Query size estimation by adaptive sampling (extended abstract)



Richard J. Lipton, Jeffrey F. Naughton

April 1990 **Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems PODS '90**

Publisher: ACM Press

Full text available: pdf(773.00 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present an adaptive, random sampling algorithm for estimating the size of general queries. The algorithm can be used for any query Q over a database D such that 1) for some n , the answer to Q can be partitioned into n disjoint subsets Q_1, Q_2, \dots, Q_n , and 2) for $1 \leq i \dots$

12 Industrial papers: query processing: Native Xquery processing in oracle XMLDB



Zhen Hua Liu, Muralidhar Krishnaprasad, Vikas Arora

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data SIGMOD '05**

Publisher: ACM Press

Full text available: pdf(253.13 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

With XQuery becoming the standard language for querying XML, and the relational SQL platform being recognized as an important platform to store and process XML, the SQL/XML standard is integrating XML query capability into the SQL system by introducing new SQL functions and constructs such as XMLQuery() and XMLTable. This paper discusses the Oracle XMLDB XQuery architecture for supporting XQuery in the Oracle ORDBMS kernel which has the XQuery processing tightly integrated with the SQL/XML engine ...

13 TinyDB: an acquisitional query processing system for sensor networks



Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong

March 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 1

Publisher: ACM Press

Full text available: pdf(1.67 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (*sampled*) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for controlling ...

Keywords: Query processing, data acquisition, sensor networks